

Knowledge-Driven Adaptive Execution of Care Pathways Based on Continuous Planning Techniques*

Gonzalo Milla-Millán¹, Juan Fdez-Olivares¹, Inmaculada Sánchez-Garzón¹,
David Prior¹, and Luis Castillo²

¹ Dpto. de Ciencias de la Computación e I.A., Universidad de Granada, Spain
{gmillamillan, faro, isanchez, dprior}@decsai.ugr.es

² IActive Intelligent Solutions, Granada, Spain
LCastillo@iactive.es

Abstract. Care pathways are excellent tools for the standardization of care delivery and the improvement of clinical efficiency. The high dynamism and unpredictability of the clinical environment require pathways to be adaptable to the deviations that may arise during their execution. In this work we present a methodology for the identification, monitoring, detection and managing of these deviations. Such deviations include evolving patient conditions, arbitrary medical modifications and unpredicted clinical settings. The care pathways are dynamically generated based on a knowledge-driven planning process that personalizes treatments according to up-to-date patient conditions and guarantees adherence to clinical guidelines recommendations. The implementation of the proposed methodology in a domain-independent continuous planning architecture is presented.

Keywords: Personalization and adaptation of healthcare processes, Managing flexibility and exceptions in healthcare processes, Lifecycle management for healthcare processes, Process modeling in healthcare, Compliance of healthcare processes

1 Motivation

Clinical guidelines and care pathways are clinical tools used in Clinical Decision Making in order to unify criteria according to the best scientific evidence[9,19]. Care Pathways can be seen as a way for adapting clinical guidelines to local conditions and their subsequent application to clinical practice[4]. These pathways are patient-centric processes that manage the patient care through the centralization of multidisciplinary tasks[13] which (1) are temporally annotated, (2)

* This work has been partially supported by the Spanish MICIIN projects TIN2008-06701-C03 and TIN2011-27652-C03-03 and the Andalusian Regional Ministry of Innovation under project P08-TIC-3572

may require the involvement of clinical resources, and (3) are intended to cover the whole patient care journey.

Thus, care pathways are an excellent tool for the patients to know how his/her care treatment is expected to evolve, for the clinical staff to have their tasks properly scheduled and for institutions to efficiently use the clinical resources. However, the clinical environment is highly dynamic and unpredictable so a key problem is how to adapt care pathways to the deviations that may occur during their execution, guaranteeing that the adapted pathways also maintain adherence to guideline recommendations.

In order to cover these needs we present an approach for the execution, monitoring and adaptation of care pathways to the evolving patient and clinical conditions, maintaining the adherence to guideline recommendations. In our proposal care pathways are dynamically generated, following a knowledge-driven process based on temporal Hierarchical Task Network (HTN) planning[5]. This process takes as inputs both an HTN-based formal representation of a clinical guideline and data extracted from a Virtual Medical Record (VMR) which includes patient and resource information. This HTN formalism has already been proved to be an enabling technology for representing clinical guidelines and generating patient-tailored care pathways from them[6,7]. It takes into account the specific patient features and observation values for generating a patient-centric and individualized care pathway, which is one of the challenges to be confronted in the Healthcare area since each patient constitutes an unique case[2]. These care pathways take the form of temporally and semantically annotated plans.

In this work we present how these plans can be executed following a continuous planning model which (1) automatically identifies the potential sources of deviation from the planned pathway, (2) monitors them and (3) adapts the pathway to the new conditions in case that any deviation from the expected course is detected. The pathway adaptation is done via a replanning process based on the same temporal HTN planning process, which re-generates the pathway according to both the knowledge encoded in the formal guideline and the current health conditions of the patient which are continuously updated in the VMR during pathway execution. This way, adherence of adapted pathways to guideline recommendations is guaranteed, as well as the observation of up-to-date clinical information.

The proposed execution model has been implemented in a planning architecture called PELEA[1] which follows a continuous planning approach. The implementation contemplates the use of a VMR which acts as an intermediate gateway that integrates patient data, clinical conditions and available resources coming from external Electronic Health Records (EHR) or Hospital Information System (HIS). Full access to this VMR is ensured at any stage of the care pathway execution management.

The implementation in PELEA is aimed at carrying out an experimentation focused on testing our methodology and demonstrating a way of implementing guidelines in clinical practice through the use of patient-tailored care pathways extracted from them. The following aspects concerning the management of clin-

ical pathways can be considered as contributions: (1) the representation of care pathways as temporal plans allows for their dynamic generation and adaptive execution; and (2) plan monitoring and execution techniques are valid to adapt care pathways to the deviations that may arise during execution. This work shows also how adherence of the adapted pathways to evidence-based guidelines is guaranteed by the planning process, which takes as input a formal representation of a clinical guideline. Moreover, all the care pathways within this approach are individualized for specific patients and tailored to available resources, thanks to the use of a VMR which gathers information about patient data and clinical conditions, as well as the clinical settings where care pathways are executed.

In the rest of the paper we first explain the knowledge representation model and how it is used to dynamically generate and adaptively execute individualized care pathways adhering to guidelines recommendations. Then, we explain our methods for the adaptive execution of the pathways. Afterwards, we explain the PELEA architecture and how these methods have been implemented there. The experiment follows together with a discussion about the results. The Related Work and the Conclusions are exposed at the end.

2 Knowledge Representation for the Dynamic Generation and Adaptive Execution of Care Pathways

The knowledge representation includes both a context model devoted to model and represent clinical data and an expert knowledge model intended to encode the clinical decisions and procedures of a clinical guideline.

The context model follows the standard recommendations of UML (Unified Modeling Language [3]) and includes those objects, properties and relations involved in a clinical guideline like patient, drugs, staff, mechanical resources, etc.

The expert knowledge model follows an extended Hierarchical Task Networks (HTN) representation which incorporates a graphical notation[18] inspired by Business Process Management (BPM) standards[22]. Its basic concepts are: *Compound Task* (or *Goal*), *Decomposition Method* and *Primitive Task*. Compound tasks represent high-level processes or goals to be accomplished in possibly alternative ways depending on the context. Decomposition methods are used to specify the alternative ways in which either a high-level process can be decomposed or a high-level goal can be achieved. Each method has related *applicability conditions* and *subtasks*. Applicability conditions are logical expressions referring to properties and relations of the context model. If applicability conditions for a method m of a task t hold in a given context, then subtasks of m specify the way to decompose (or accomplish) the compound task (or goal) represented by t in that given context. Tasks are decomposed via their decomposition methods until all the activities to be performed are primitive tasks. Primitive tasks represent concrete actions which involve a change in the context data when executed. The relations between tasks, decomposition methods and primitive tasks allow the representation of the clinical decisions and procedures of a clinical guideline in terms of a compositional hierarchy of tasks.

The dynamic generation of patient-tailored care pathways is based on a knowledge-driven, temporal hierarchical planning process described in [6,7]. This planning process receives as inputs the expert knowledge model which encodes guideline recommendations, a starting date for the pathway and context data accomplishing the UML specification of the context model. As explained in [6], the temporal HTN planning process allows the representation and reasoning about temporal constraints derived from hierarchical decompositions is able to express different kinds of periodic patterns to be followed by temporal constraints. The so generated care pathways are sets of instantiated primitive tasks that adhere to guideline recommendations. These tasks are annotated with information about (1) their start and end time points, (2) the changes that they produce in the context data, and (3) the conditions that must hold during their execution.

2.1 Use of the Knowledge Representation for the Adaptive Execution of Pathways

The knowledge representation can be used to represent the management processes for several kind of deviations that may arise during pathway execution. Following the definitions of [20], all the deviations that we are able to detect and manage are *predictable*, in the sense that they can be foreseen at design time.

The first two potential sources of deviation are *decision points* (DP) and *unplanned patient conditions* (UPC). Both of them can be found in clinical guidelines recommendations. DPs represent choices that must be done among different treatment alternatives depending on patient conditions at a given instant. UPC refer to patient conditions which require additional care tasks and may arise (or not) during pathway execution. The difference among them is that DPs are precisely located in time, while it is not exactly known when UPCs will arise, if ever. The management of both kinds of deviations is encoded as different decomposition methods in the expert knowledge model, but the monitoring and detection of deviations is done in an unified way as will be seen in next section. Figure 1(a) models a compound task for DP management. Applicability conditions of each method relate to patient conditions and subtasks of each method model the different treatment alternatives. Figure 1(b) models a compound task for UPC management. Applicability conditions of its two decomposition methods distinguish whether UPC are present or not. Just in case, the additional care tasks are applied. When any of these compound tasks is included in the expert knowledge model, the planning process must know which method to apply, which ultimately depends on patient conditions. For that reason, the *expected patient evolution* must be encoded as context data, assuming (1) more foreseeable patient conditions at DPs and (2) no-appearance of UPCs. These expected patient evolution is encoded as patient attributes in the context model and as changes in the context data made by specific primitive tasks. This way the planning process can take this expected patient evolution into account and the generated care pathways can be temporally annotated with this information. During pathway execution, we will need to monitor at proper time if real patient evolution matches the expected one.

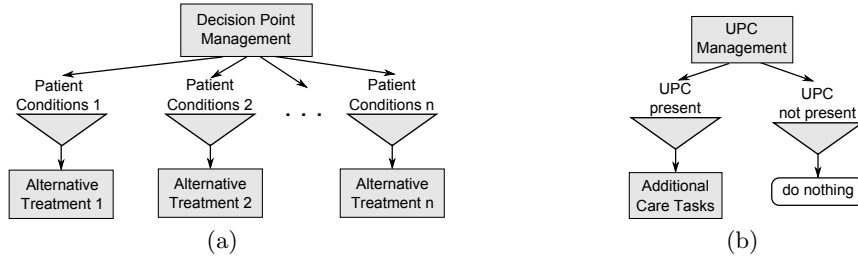


Fig. 1. Modelling of Decision Points and UPC management

Another deviation type is the practitioner request for delaying the pending treatment. This is an example of arbitrary medical modification. The planning process takes as input the starting date for the pathway, which can be used to manage this kind of deviation. During pathway execution, we will need to continuously monitor this potential practitioner request.

Last source of deviation is resource availability. Committed resources may be unavailable during pathway execution due to several reasons. The resource availability can be encoded as object attributes in the UML-based context model which can be dynamically changed by those primitive tasks where they are involved. During pathway execution we will need to monitor availability for each committed resource at each time they are scheduled.

Summing up, the context model is used as a means for representing the expected patient evolution and the resource availability. Compound tasks and decomposition methods allow for the modeling of management processes for decision points and unplanned patient conditions. Primitive tasks are the way to model the changes to be done in the context data, which include patient conditions and resource availability. Furthermore, care pathways are annotated with the changes to be made to the context data together with the time at which these changes must be made. Finally, the delay of the pending treatment can be managed with the starting date of the pathway.

3 Adaptive Execution of Care Pathways Based on Continuous Planning Techniques

In this section, we present our methodology for the adaptive execution of care pathways based on continuous planning techniques. All the deviations that can be managed are *predictable* as defined in [20], in the sense that they can be foreseen at design time. Thus, all the management processes for these deviations are encoded in the expert knowledge model as explained in the previous section. However, the adaptive process requires several stages which make use of this knowledge. Figure 2 and algorithm 1 depict the main elements and overall process.

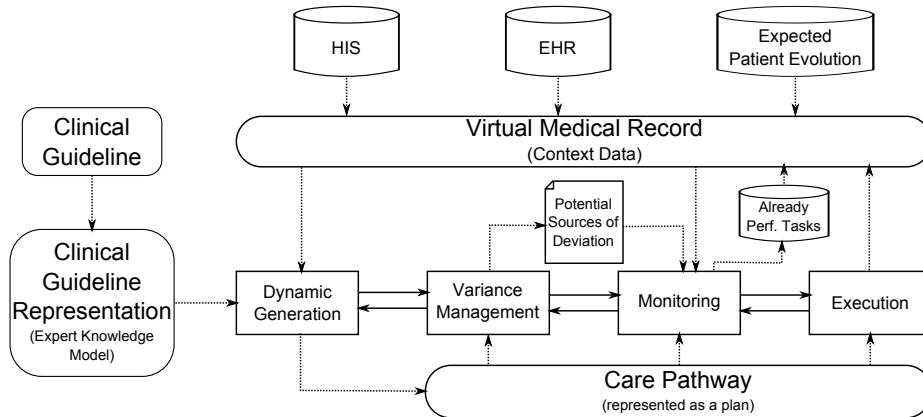


Fig. 2. Adaptive Execution of Care Pathways Based on Continuous Planning

The main elements in our approach are a formal representation of the Clinical Guideline, a VMR, the Care Pathway (which takes the form of a temporally annotated plan) and its potential sources of deviation during execution. The Clinical Guideline representation is encoded according to the expert knowledge model. The VMR implements the context data according to the UML-based context model and acts as an intermediate gateway that integrates patient data, clinical conditions and available resources coming from external EHR and/or HIS, as well as the expected patient evolution. Moreover, information about already performed care tasks is also embedded in the VMR, to avoid their repetition when adapting a pathway, as we will see next. The extraction of a patient-tailored Care Pathway from these elements following an HTN planning process has already been explained in previous works[6,7]. In this sense, our work extends that approach with adaptive execution capabilities. In order to endow our system with these execution capabilities we need also to know the potential sources of deviation from pathway. Each of these potential sources of deviation consists of (i) a condition that need to be monitored in the VMR data, (ii) the time when this condition is expected to hold. The system considers that a deviation arises every time that one (or more) of these conditions do not hold at the expected time.

The main processes of our methodology are: (1) a *Dynamic Generation* process, which generates a plan representing the care pathway from the clinical guideline formal representation and the VMR; (2) the *Variance Management* process, the goal of which is to extract the potential sources of deviation to be monitored during pathway execution; (3) a *Monitoring* process in charge of both sending actions to execute and checking if any deviation arises; and (4) the *Execution* of the care tasks which continuously update the VMR accordingly to primitive tasks specifications. Full access to the VMR is ensured at any stage of the care pathway execution management.

Algorithm 1: Care Pathway Dynamic Generation and Adaptive Execution

Input: Virtual Medical Record(*VMR*), Clinical Guideline representation in HTN form(*Guide*)

- 1 Set already performed tasks (*APT*) = \emptyset ;
- 2 Get actual patient state (*PatSta*), expected evolution (*ExpEvo*), resource information (*ResInf*) and already performed tasks (*APT*) from the *VMR*;
- 3 Generate Care Pathway (*Pathway*) from *Guide* applying *PatSta*, *ExpEvo*, *ResInf* and *APT*;
- 4 Wait expert approval for *Pathway*;
- 5 Compute potential sources of deviation (*PotDev*) from *Pathway*;
- 6 **repeat** /* Execution-Monitoring stage */
- 7 Execute next action(s) (*NexAct*) of *Pathway* according to *start-time* and update *VMR* accordingly;
- 8 **foreach** action in *NexAct* (ordered by end-time) **do**
- 9 **if** any deviation of *PotDev* at end-time of action is detected **then**
- 10 | Go back to step 2;
- 11 **else** /* action was performed successfully */
- 12 | Store action in *APT*;
- 13 **end**
- 14 **end**
- 15 **until** No remaining actions to execute in *Pathway*;

The full life-cycle of the methodology is as follows. First, (lines 1-3 of algorithm 1) *Dynamic Generation* generates a care pathway according to guideline recommendations taking as input the actual patient conditions and clinical settings from the *VMR*, in order to individualize the pathway and tailor it to available resources. To be able to plan for the long-term also the expected patient evolution is gathered from the *VMR*. Initially, no care tasks have been applied to the patient so information about already performed tasks is empty. The generated care pathway takes the form of a plan and is annotated with temporal information about (1) *start-time* and *end-time* points for each care task³, (2) expected evolution of the patient, (3) unplanned patient conditions that may arise according to guideline specifications, and (4) resources availability through the pathway execution. From there, the *Variance Management* process (line 5) computes the potential sources of deviation that need to be monitored during pathway execution. Every potential source of deviation consists of (i) a condition to be monitored in the *VMR* data, (ii) the time when this condition must hold. This data is extracted from the effects of the primitive tasks in the care pathway, so a potential source of deviation is extracted for every effect of every primitive task and is related to the time that the effect is expected to hold. For instance, $\{(Patient.good\text{-}level\text{-}neutrophils\ Carl), 10\}$ is a potential source of deviation which indicates that patient *Carl* must have a good level of neutrophils at tenth day after the beginning of pathway execution, otherwise it is considered

³ Imprecise times are not allowed but rich and complex time constraints can be encoded as explained in [6].

a deviation that must be managed. After these potential sources of deviation have been computed, the *Execution-Monitoring* stage (lines 6-15) starts. There, the *Monitoring* process continuously checks for the actions that need to be executed (according to their start-time) and send them to execution. The next set of actions to be executed is computed in a way so there is no overlapping of any action in the set with any other action of a previous or future set. For every executed action, the VMR is updated accordingly (line 7). Then, *Monitoring* checks if any deviation has arisen at that moment (line 9). If everything goes as expected it updates the already performed tasks with the successfully applied actions (line 12) and continues sending actions to execution. If a deviation from the expected pathway is detected (i.e. the level of neutrophils of patient Carl is not good at tenth day after the beginning of the pathway execution), then *Monitoring* notifies it to *Variance Management* which is in charge of calling again the *Dynamic Generation* process (line 10). *Dynamic Generation* will get the up-to-date information from the VMR - now the information about the already performed tasks will not be empty - together with the guideline recommendations from the formal guideline representation and generates a new plan representing the new pathway. This new pathway covers the whole patient care journey from the moment the deviation was detected till the end of the treatment. The new pathway is obtained via a replanning process, based on the same temporal HTN planning process, which re-generates the pathway according to both the knowledge encoded in the formal guideline - which encodes also the management processes for the deviations - and the current health conditions of the patient, continuously updated in the VMR during pathway execution. The cycle is repeated until no more actions need to be executed (e.g. no more care tasks to apply) (line 15). All the generated pathways must be approved by the person responsible (line 4).

4 Implementation in a Continuous Planning Architecture

The explained methodology has been implemented in a modular architecture called PELEA[1] which follows a continuous planning approach. PELEA allows the rapid prototyping of planning applications using standard planning techniques. Figure 3 depicts the architecture, where the gray shadow surrounds the modules that we have made use of. For a more detailed explanation of the architecture we refer the reader to [1].

The *high-level replanner* module of PELEA implements the *Dynamic Generation* process. An HTN planner is in charge of obtaining a temporally annotated plan which represents the care pathway. This pathway representation is obtained taking the formal guideline representation and the VMR data, following a temporal HTN planning process. Furthermore, the pathway representation is annotated with the required information to compute the potential sources of deviation. The *high-level replanner* sends this output to the *decision support* module. The *decision support* module of PELEA is in charge of computing the conditions that must be monitored during execution and implements the *Vari-*

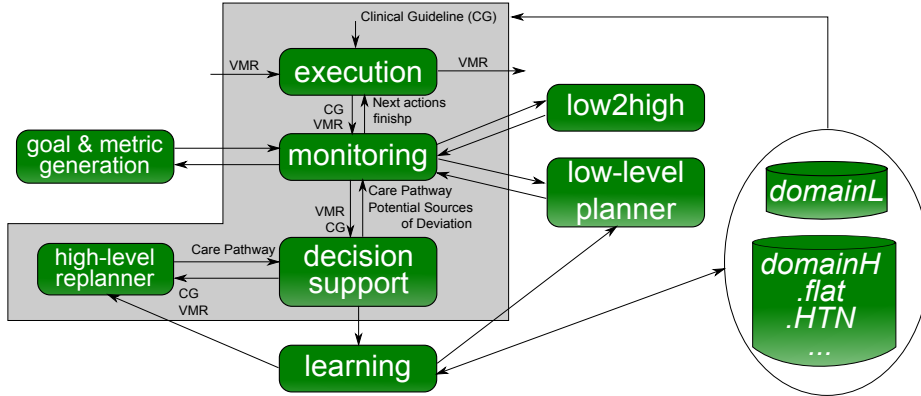


Fig. 3. PELEA Architecture

ance Management process, which computes potential sources of deviation in the following way: (1) it sorts the primitive tasks of the pathway representation by their end-time, (2) for every effect (change to context data) of every primitive task, it adds a potential source of deviation, (3) each of these potential sources of deviation consists of a condition that is expected to hold (the expected effect of the primitive task) and the time when it is expected to hold (the end-time of the primitive task). The care pathway and these potential sources of deviation are sent by the *decision support* to the *monitoring* module. The *monitoring* module of PELEA is in charge of checking the effects of every primitive task at its end-time and has been used to implement the *Monitoring* process. The *monitoring* module is in charge of sending next actions to *execution* module, according to their start-time. The *execution* module of PELEA is in charge of executing actions and updating the context data accordingly. It implements the *Execution* process which represents the real clinical settings where the system would be implemented and therefore, actions would be executed. The *execution* module receives the context data (i.e. the VMR) and the formal clinical guideline representation, making them accessible to the rest of modules through the data flow, guaranteeing full access to VMR at any stage of the care pathway execution.

5 Experiment

The goal of this experiment is to check that our methodology is able to detect and manage deviations from pathway during execution, adapting these pathways in a way that adherence to guidelines and use of up-to-date clinical information is guaranteed. For this purpose, we simulate these deviations using a probabilistic simulator as *Execution* process (see figure 2), embedded in the *execution* module of PELEA architecture (see figure 3). The probabilistic simulator uses

a modified version of the actions (primitive tasks used in the formal guideline to represent care tasks) to simulate their execution. Thus, their *expected* effects (changes that they made on the context data) are assigned a probability. Then, several (independent) sets of effects with related probabilities are added to each action. The probabilities distribution has been manually created. To execute an action, the simulator selects one of the related sets of effects following the related probability distribution. The monitoring process always checks for the *expected* effects after each action is executed. Therefore, if the simulator chooses one of the new added sets of effects, it is considered by the monitoring process as a deviation from the expected pathway.

A simplification of a guideline for planning the treatment of Hodgkin’s disease [10], as example of clinical guideline, has been used for the experiment.

In this guideline there are several decision points related to evaluations of tumor remission in the patient, being *complete-remission* and *partial-remission* the possible values of these evaluations. For the experiment, we assume that the expected evolution of the patient is *complete-remission* and specify it in the VMR as a patient attribute. In the expert knowledge model, a primitive task for the evaluation is modeled with associated effect (*Patient.complete-remission ?patient*). For the probabilistic simulator, the corresponding action will have two different possible effects: (1) (*Patient.complete-remission ?patient*) and (2) (*Patient.partial-remission ?patient*) with a related probability distribution. This way, the *execution* module is able to simulate the expected value (*complete-remission*) but also (*partial-remission*), which should be detected as a deviation by *monitoring*. For patients of group 1 - which is the case that we present below - just a decision point of this kind is specified in [10], for which the management process for deviations (e.g. a *partial-remission* of the tumor) is to apply 14 radiotherapy sessions. According to the definitions in [20], those deviations coming from decision points are synchronous. In this experiment they are treated in this way, detecting them at the same time that they occur, if ever.

On the other hand, *bad level of platelets*, *bad level of neutrophils* and *fever* have been modeled as unplanned patient conditions. The expected patient evolution in the context data assumes that the patient will not show any of these circumstances during the pathway execution. To detect this potential deviation we have set the conjunction of (*Patient.good-level-platelets ?patient*), (*Patient.good-level-neutrophils ?patient*), and (*not (Patient.has-fever ?patient)*) as the effects of the primitive task which represents a drug administration. Thus, a potential source of deviation is detected for each one of these facts after each drug administration. Furthermore, we add the rest of possible combinations of conditions (i.e. a bad level of neutrophils and/or bad level of platelets and/or fever) to the related action in the probabilistic domain, with a related distribution probability. This way, the *execution* module will be able to simulate both the expected values and different kind of deviations. The management processes specified in the formal guideline for these unplanned patient conditions are: (1) if a *bad level of platelets* is detected, perform a blood transfusion; (2) if a *bad level of neutrophils* is detected, apply ten Filgrastim administrations and one *increase-neutrophils*

special action⁴; and (3) if the patient has *fever*, apply four administrations of Cotrimoxazol and one *reduce-fever* special action⁵. Unplanned patient conditions originate asynchronous deviations which in this experiment are monitored at discrete instants of time (related to the end of every drug administration), being possible that there is a delay from the moment the deviation occurs and its detection by the system. The improvement of this feature is planned as future work by the authors.

To simulate a possible practitioner request for delaying the pending treatment, the primitive task for radiotherapy sessions is modeled with (*no-delay*) as expected effect, so these requests are monitored after each radiotherapy session⁶. The corresponding action in the probabilistic domain has both effects of (*no-delay*) and (*delay*) with a related probability distribution. If the simulator chooses the *delay* effect, the *monitoring* process will detect a deviation due to a practitioner request for delaying the pending treatment. The beginning-of-treatment date is used instead of the formal guideline to manage this kind of deviation which, like UPCs, is asynchronous but monitored at discrete instants of time (related to the end of every radiotherapy session).

Resource unavailability has not been worked in this preliminary experiment. However, we explain how they are correctly rescheduled for each care pathway adaptation, since the resource information is stored in the VMR.

5.1 Results

In this preliminary experiment we have extracted a patient-tailored care pathway for a patient of gender *male* and *group 1*, according to the specifications of [10]. The initial care pathway was composed of 131 care tasks. Its execution in PELEA lasted 63 minutes and 11 seconds and 20 deviations were thrown by the probabilistic simulator. All the deviations were detected and correctly managed with the subsequent adaptation of the pathway. The total number of executed care tasks was 193. Table 1 shows a summary of the deviations, their type, number, what the adaptation of the pathway consisted of and the number of actions that was added in each case.

The results of this experiment show that every simulated deviation is detected and managed. For each deviation, the correct adaptation is done by the *Dynamic Generation* process, tailoring a re-generated pathway which adheres to guideline recommendations and takes into account the new patient conditions. Resources are also rescheduled when necessary in this re-generation process.

⁴ This special action is intended to represent that the intention of applying ten Filgrastim administrations is to increase the level of neutrophils.

⁵ This special action is intended to represent that the intention of applying four Cotrimoxazol administrations is to reduce the fever.

⁶ Though this should be checked continuously we have chosen to do it only after this primitive task for the first approximation of the experimentation.

Table 1. Experiment Results. Line 1, for instance, is read: “A partial-remission of tumor was simulated once. Care pathway was adapted with 14 additional radiotherapy sessions, which is the proper alternative treatment encoded in the formal guideline.”.

	Deviation Type	#	Adaptation	# added actions	
Decision Points	partial-remission	1	14 radioth. sessions	14	
Unplanned Patient Conditions	bad level platelets	12	1 blood transfusion per each	12	
	bad level neutrophils	1	10 Filgrastim admins. + 1 increase-neutrophils	11	
	fever	5	(4 Cotrimoxazol admins. + 1 reduce-fever) per each	25	
Medical Intervention	delay of treatment	1	treatment is delayed and resources rescheduled	0	
		Total	20	Total	62

6 Related Work

Previous works of our group shows that HTN planning is an enabling technology for the representation of clinical guidelines and the generation of treatment plans [6], and also how Computer Interpretable Guidelines (CIGs) can be translated into HTN planning domains so patient-tailored care pathways can be extracted from them [7]. The work here presented adds adaptive execution capabilities for the care pathways. Another recent work [21] is focused on the medical exceptions representation and the differentiation of normal and exceptional flow in the formal guideline modeling.

There is an intensive review of computer-based execution of clinical guidelines in [12]. Though we set the focus on the execution of care pathways, there are common points because of the guideline adherence of the patient-tailored care pathways here treated. In this sense, our approach has all the elements needed to encode a guideline [6,7] and adds execution capabilities to care pathways that adhere to so-encoded guidelines.

In [14] a multiagent approach is used to manage test requests for patients. The scheduling of the test requests involves resource management and can be seen as a plan that is monitored and adapted in case that unexpected events occur or a high priority request needs of already scheduled resources. However, these unexpected events are not clearly stated and no care tasks are observed, so it is not clear how this approach would behave for other critical aspects as the ability to compute the proper dose of a drug administration depending on patient features. It is something that our approach can deal with [6].

Quite close to ours are the works [15,16]. We use a different terminology but the differences go beyond. These works are focused on high-frequency domains and short term plans, while the one presented here is frequency independent and more focused on the long term. The monitoring in these approaches is done in fixed time intervals while we check only those times where a deviation can imply a change to the pathway. Moreover, these approaches face the plan generation

with a non-deliberative process, based on the static adaptation of skeletal-plans, that is not interleaved with temporal constraints reasoning as in ours.

Interesting works based on workflows can be found in the literature [17,8]. They use ECA rules for the local modification of workflows that which represent clinical guidelines. They do not provide the complete set of expected care tasks as we do, which allows for a better and earlier resources scheduling, as well as the awareness of the personnel involved. Moreover, our approach supports the temporal dimension of failures, solving the problem cited in [17].

The work of [11] makes the pathway adaptation based on ratings given by physicians which is more time consuming and biased than guaranteeing standard guidelines adherence as in our approach. In [2] dynamic adaptation occurs during each step of care pathway execution by applying semantic rules that handle the deviations, but no explicit reference to guideline adherence is given and the care pathways are not extracted for the whole patient journey.

7 Conclusions

We have presented an approach based on continuous planning techniques for the dynamic generation and adaptive execution of individualized care pathways which adhere to evidence-based guidelines recommendations. This approach gives the expected course of events for the particular situation of the patient in the form of a care pathway represented as a plan. It is also aware of the potential sources of deviation from this normal course and is able to manage those predictable complications that may arise, adapting the normal course to the new circumstances. The knowledge representation model allows the representation of (1) clinical guidelines, (2) context data, (3) temporally annotated care pathways, and (4) management processes for several kind of deviations. The implementation of the approach in a continuous planning architecture has been presented. A preliminary experiment shows that our approach is able to dynamically adapt pathways to deviations that may arise during their execution. Guideline adherence of the adapted pathways is guaranteed by the underlying planning-based generation process. Full access to the context data at every stage of the pathway execution ensures the use of continuously up-to-date information. Moreover, adapted pathways are stored together with the context data at the moment they were generated so it can serve for statistical analysis of deviations.

References

1. V. Alcázar, C. Guzmán, D. Prior, D. Borrajo, L. Castillo, and E. Onaindía. Pelea: Planning, learning and execution architecture. In *UK Planning and Scheduling Special Interest Group (PlanSIG)*, 2010.
2. D.A. Alexandrou, I.E. Skitsas, and G.N. Mentzas. A holistic environment for the design and execution of self-adaptive clinical pathways. In *Information Technology and Applications in Biomedicine (ITAB)*, 2009.
3. G. Booch, J. Rumbaugh, and I. Jacobson. *The unified modeling language user guide*. Inc., 1999.

4. Harry Campbell, Rona Hotchkiss, Nicola Bradshaw, and Mary Porteous. Integrated care pathways. *BMJ*, 316(7125):133–137, 1 1998.
5. L. Castillo, J. Fdez-Olivares, O. García-Pérez, and F. Palao. Efficiently handling temporal knowledge in an htn planner. In *International Conference on Automated Planning and Scheduling (ICAPS)*, 2006.
6. J. Fdez-Olivares, L. Castillo, J.A. Cózar, and Ó. García Pérez. Supporting clinical processes and decisions by hierarchical planning and scheduling. *Computational Intelligence*, 27:103–122, 2011.
7. A. Gonzalez-Ferrer, A. ten Teije, J. Fdez-Olivares, and K. Milian. Automated generation of patient-tailored electronic care pathways by translating computer-interpretable guidelines into hierarchical task networks. *Artificial Intelligence In Medicine*, 2012.
8. U. Greiner, R. Mueller, E. Rahm, J. Ramsch, B. Heller, and M. Loeffler. Adaptflow: protocol-based medical treatment using adaptive workflows. *Methods of Information in Medicine*, 44:80–88, 2005.
9. J.M. Grimshaw and I.T. Russell. Effect of clinical guidelines on medical practice: a systematic review of rigorous evaluations. *The Lancet*, 342:1317 – 1322, 1993.
10. SEOP Working Group. National protocol for diagnosis and treatment of hodgkins disease - eh-seop.003. Technical report, SEOP, May 2000. (in Spanish).
11. Z. Huang, X. Lu, and H. Duan. Using recommendation to support adaptive clinical pathways. *Journal of Medical Systems*, 36:1849–1860, 2012.
12. David Isern and Antonio Moreno. Computer-based execution of clinical guidelines: A review. *International Journal of Medical Informatics*, 77(12):787 – 808, 2008.
13. Steve Kelsey. Managing patient care: are pathways working? *Practice Development in Health Care*, 4(1):50–55, 2005.
14. Catherine C Marinagi, Constantine D Spyropoulos, Christos Papatheodorou, and Stavros Kokkotos. Continual planning and scheduling for managing patient tests in hospital laboratories. *Artificial Intelligence in Medicine*, 20(2):139 – 154, 2000.
|ce:title|Planning and Scheduling in the Hospital|ce:title|.
15. Silvia Miksch, Werner Horn, Christian Popow, and Franz Paky. Vie-vent: knowledge-based monitoring and therapy planning of the artificial ventilation of newborn infants. *Artificial Intelligence in Medicine*, 10:218–229, 1993.
16. Silvia Miksch and Andreas Seyfang. Continual planning with time-oriented, skeletal plans. In *European Conference on Artificial Intelligence (ECAI 2000)*, pages 511–515, 2000.
17. Robert Müller, Ulrike Greiner, and Erhard Rahm. Agentwork: a workflow system supporting rule-based workflow adaptation. *Data & Knowledge Engineering*, 51(2):223 – 256, 2004.
18. F. Palao, J. Fdez-Olivares, L. Castillo, and O. García. An extended htn knowledge representation based on a graphical notation. In *Workshop Knowledge Engineering for Planning and Scheduling, ICAPS*, 2011.
19. M. Panella, S. Marchisio, and F. Di Stanislao. Reducing clinical variations with clinical pathways: do pathways work? *International Journal for Quality in Health Care*, 15:509–521, 2003.
20. Mor Peleg, Judith Somekh, and Dov Dori. A methodology for eliciting and modeling exceptions. *Journal of Biomedical Informatics*, 42(4):736 – 747, 2009.
21. Inmaculada Sánchez-Garzón, Juan Fernández-Olivares, and Luis Castillo. An approach for representing and managing medical exceptions in care pathways based on temporal hierarchical planning techniques. In *ProHealth'12 / KR4HC'12*, 2012.
22. S White. Introduction to bpmn. Technical report, IBM Cooperation, 2004.